

**BAYESIAN TREATY MONITORING: PRELIMINARY REPORT**

Stuart J. Russell,<sup>1</sup> Stephen C. Myers,<sup>2</sup> Nimar S. Arora,<sup>1</sup> David A. Moore,<sup>1</sup> and Erik Sudderth<sup>3</sup>

University of California, Berkeley<sup>1</sup>, Lawrence Livermore National Laboratory<sup>2</sup>, and Brown University<sup>3</sup>

Sponsored by the Defense Threat Reduction Agency

Award Nos. BRBAA08-0042  
Proposal No. HDTRA1-11-1-0026

**ABSTRACT**

Our project has initiated and will develop and evaluate a new Bayesian approach for nuclear test monitoring. We anticipate that the new approach will yield substantially lower detection thresholds, possibly approaching a theoretical lower bound that we hope to establish. We will also develop new techniques to implement such monitoring capabilities within a general-purpose Bayesian modeling and inference system that may eventually support a wide range of information-system needs for arms treaties.

In ongoing work that is moving towards possible deployment, we have completed a prototype seismic monitoring system based on a generative, vertically integrated statistical model linking hypothesized events to “detections” extracted from raw signal data by commonly used algorithms. On test data sets of naturally occurring events curated by human experts, our system exhibits roughly 60% fewer detection failures than the currently deployed automated system, SEL3, that forms part of the International Monitoring System.

The current phase of the project moves away from hard-threshold detections altogether. Instead, the generative model spans the full range from events to measured signal properties. Given the observed signal traces, the statistical inference algorithm attempts to maximize a whole-network statistical measure of the likelihood that an event – or collection of events – has occurred. Specialized techniques such as *waveform matching* and *double differencing* are realized within our framework as special cases of probabilistic inference; our initial experiments using 2D simulated data indicate that a full Bayesian analysis can provide more accurate absolute and relative locations than double differencing, while simultaneously estimating the velocity structure of the observed region.

As we move toward a full-scale implementation, the primary tasks will involve the development of accurate predictive models of waveform properties. These models will combine both parametric forms (for example, triangular envelopes in multiple frequency bands) and nonparametric forms based on previously observed waveforms from nearby events. Hybrid models will smoothly interpolate between these two forms depending on the distance of the hypothesized event from previously observed events.

<b>Report Documentation Page</b>		<i>Form Approved OMB No. 0704-0188</i>
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>		
1. REPORT DATE <b>SEP 2011</b>	2. REPORT TYPE	3. DATES COVERED <b>00-00-2011 to 00-00-2011</b>
4. TITLE AND SUBTITLE <b>Bayesian Treaty Monitoring: Preliminary Report</b>		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>University of California, Berkeley, Berkeley, CA, 94720</b>		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>		
13. SUPPLEMENTARY NOTES <b>Published in the Proceedings of the 2011 Monitoring Research Review - Ground-Based Nuclear Explosion Monitoring Technologies, 13-15 September 2011, Tucson, AZ. Volume I. Sponsored by the Air Force Research Laboratory (AFRL) and the National Nuclear Security Administration (NNSA). U.S. Government or Federal Rights License</b>		

**14. ABSTRACT**

**Our project has initiated and will develop and evaluate a new Bayesian approach for nuclear test monitoring.** We anticipate that the new approach will yield substantially lower detection thresholds, possibly approaching a theoretical lower bound that we hope to establish. We will also develop new techniques to implement such monitoring capabilities within a general-purpose Bayesian modeling and inference system that may eventually support a wide range of information-system needs for arms treaties. In ongoing work that is moving towards possible deployment, we have completed a prototype seismic monitoring system based on a generative, vertically integrated statistical model linking hypothesized events to ?detections? extracted from raw signal data by commonly used algorithms. On test data sets of naturally occurring events curated by human experts, our system exhibits roughly 60% fewer detection failures than the currently deployed automated system, SEL3, that forms part of the International Monitoring System. The current phase of the project moves away from hard-threshold detections altogether. Instead, the generative model spans the full range from events to measured signal properties. Given the observed signal traces, the statistical inference algorithm attempts to maximize a whole-network statistical measure of the likelihood that an event ? or collection of events ? has occurred. Specialized techniques such as waveform matching and double differencing are realized within our framework as special cases of probabilistic inference; our initial experiments using 2D simulated data indicate that a full Bayesian analysis can provide more accurate absolute and relative locations than double differencing, while simultaneously estimating the velocity structure of the observed region. As we move toward a full-scale implementation, the primary tasks will involve the development of accurate predictive models of waveform properties. These models will combine both parametric forms (for example triangular envelopes in multiple frequency bands) and nonparametric forms based on previously observed waveforms from nearby events. Hybrid models will smoothly interpolate between these two forms depending on the distance of the hypothesized event from previously observed events.

**15. SUBJECT TERMS**

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>10</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

## **OBJECTIVES**

Our objective is to realize dramatic improvements in sensitivity, accuracy, and robustness of global monitoring systems for nuclear tests through a novel Bayesian approach to whole-network data analysis.

The project will develop a new mathematical and computational approach to the analysis of the sensor data collected from global monitoring systems. The approach involves the application of rigorous Bayesian statistical analysis to the entire monitoring problem and requires the development of a complete, vertically integrated, empirically validated, generative statistical model of event occurrence, signal propagation, and sensing, as well as efficient and provably correct inference algorithms for extracting the most likely event history and/or a posterior distribution over event histories from the measured sensor data. At a fundamental level, we hope to gain a much deeper and more accurate understanding of the limits of detectability than is given by a classical per-station SNR analysis. We expect to realize the following benefits: (1) Substantially more accurate and sensitive detection and localization of events, particularly at lower magnitudes. (2) A monitoring software architecture that allows straightforward incorporation of multiple sensor modalities (including new modalities as they arise) and new and improved physics-based models such as source models and phase velocity and attenuation models. (3) Extensibility to other monitoring problems arising in treaty verification.

## **RESEARCH ACCOMPLISHED**

The period of performance for the DTRA-funded project has not yet begun at the time of writing; therefore, this report covers some background material, our relevant past work, some initial steps taken for the new project, and a brief summary of our planned activities.

### **Background: Bayesian monitoring**

The proposed research is motivated by a fundamental problem with current approaches to seismic monitoring, namely, the problem of robust detection of signals with low signal-to-noise ratio (SNR). Standard detection practice is to declare a detection when the average amplitude in a short time window increases above a predefined multiple of average amplitude in a long time window (a.k.a. short-term average/long-term average, STA/LTA). Independent algorithms then attempt to associate detections to seismic events. To avoid creating a flood of spurious events, a very high SNR threshold is typically used for detection. This approach is flawed in at least three important ways. First, it ignores the *absence* of detections in evaluating a proposed event. Second, it fails to take advantage of spatiotemporally *correlated* signals at multiple stations to detect that an event has occurred. Third, it discards details of the signal that support more accurate estimation of event properties via techniques such as waveform matching.

Our proposed approach is based on Bayesian inference, a method well-established in many areas of science that solve inverse problems including seismology itself (Duijndam, 1988ab). They have been very effective in tomographic applications (Taylor et al., 2001; Simmons et al., 2010b) and event localization (Myers et al., 2007, 2009) and have been applied to the general beamforming problem (Bell et al., 2000). Our work (Russell et al., 2009, 2010; Arora et al., 2009, 2010) is to our knowledge the first complete Bayesian monitoring system, solving both association and localization problems within a unified probability model.

In general, Bayesian inference yields a *posterior probability distribution* over a set of *hypotheses*  $X$  given some *evidence*  $Y = y$ . In seismic monitoring, a hypothesis is a collection of events occurring over space and time; the evidence consists of the raw seismic (and other) waveform signals from all sensors over the time period of interest. The inference process is based on a *probability distribution* or *model* with two components:

- The *prior probability distribution*  $P_\theta(X)$  over hypotheses; for the monitoring problem, this includes the natural seismicity distribution and a probability distribution for chemical and nuclear explosions on or near the Earth's surface. The subscript  $\theta$  represents the parameters of this distribution, which can be input as prior information and learned from empirical data.
- The *conditional probability distribution*  $P_\phi(Y=y | X)$  for the evidence given each possible hypothesis; in our case, this part of the model describes how signals propagate through the earth and how they are detected by sensors, as well as the ways in which noise signals arise. The subscript  $\phi$  represents the parameters of this distribution; again, these can be learned from empirical data, but there is also a great deal of physics-based knowledge that constrains the possible values. In seismology, the model of propagation is

often called the *forward model*, although it is important to note that in the Bayesian context it includes a quantitative characterization of uncertainty.

Bayes' rule simply multiplies these two components together to give the posterior probability distribution over the set of hypotheses, given the available evidence:

$$P(X \mid Y=y) = \alpha P_\phi(Y=y \mid X) P_\theta(X)$$

where  $\alpha$  is a normalizing constant. To the extent that the prior and conditional distributions correctly describe knowledge of seismicity, propagation, and so on, the posterior distribution represents an optimal inference from the available data. An inference algorithm may compute a single *most likely explanation* or MLE hypothesis, instead of the full posterior distribution over hypotheses. Because there are infinitely many possible hypotheses (each a set of seismic events), the calculations involved are nontrivial and require efficient inversion of the forward model.

As a side effect of the inference process, the Bayesian approach generates information that can be used to continuously adapt the model parameters  $\theta$  and  $\phi$  to better explain the data. This adaptation requires no “ground truth” (unlike supervised learning methods) and hence allows for continuous self-calibration and sensor diagnostics.

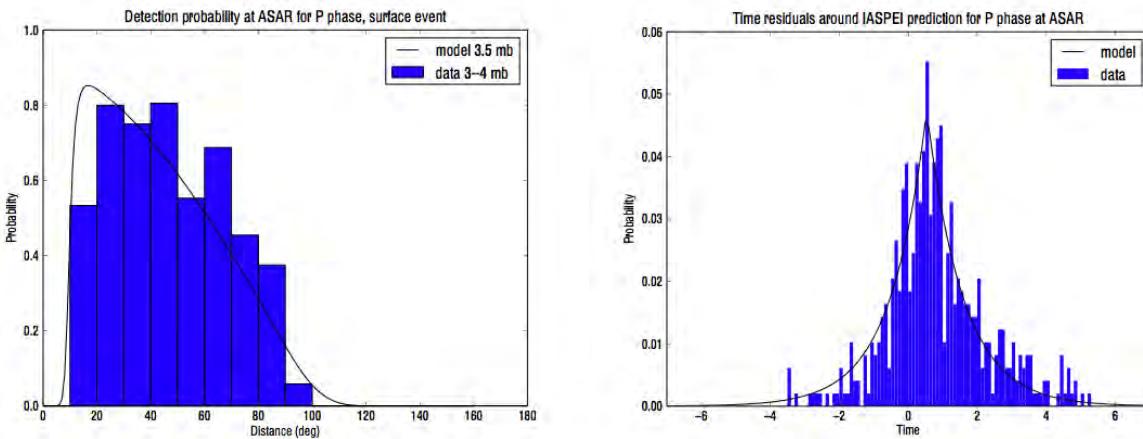
### Background: Detection-based Bayesian monitoring

When the raw data  $y$  are complicated – as in the case of complete seismic waveforms – it is common for Bayesian methods to deal with a simplified representation  $f(y)$  of the raw data. Here  $f$  can be any deterministic function. In *detection-based Bayesian monitoring*,  $f$  represents the signal processing algorithms that are used to pull out detections from the signal traces at each station. That is,  $f(y)$  is the set of all detections (with associated characteristics) from all stations during the period of interest. The performance of a detection-based Bayesian monitoring system is limited by the underlying detection algorithm  $f$ . Such an algorithm – for example, thresholding based on STA/LTA – may fail to detect some arriving signals while generating spurious detections and abstracting away important details of true signals. Nonetheless, as we show below, the performance of a Bayesian monitoring system may be significantly better than current systems that utilize the same detection algorithms.

NET-VISA (Arora et al., 2010, Arora et al., 2011) consists of a collection of probability distributions and an inference algorithm that computes an MLE hypothesis. The probability distributions include:

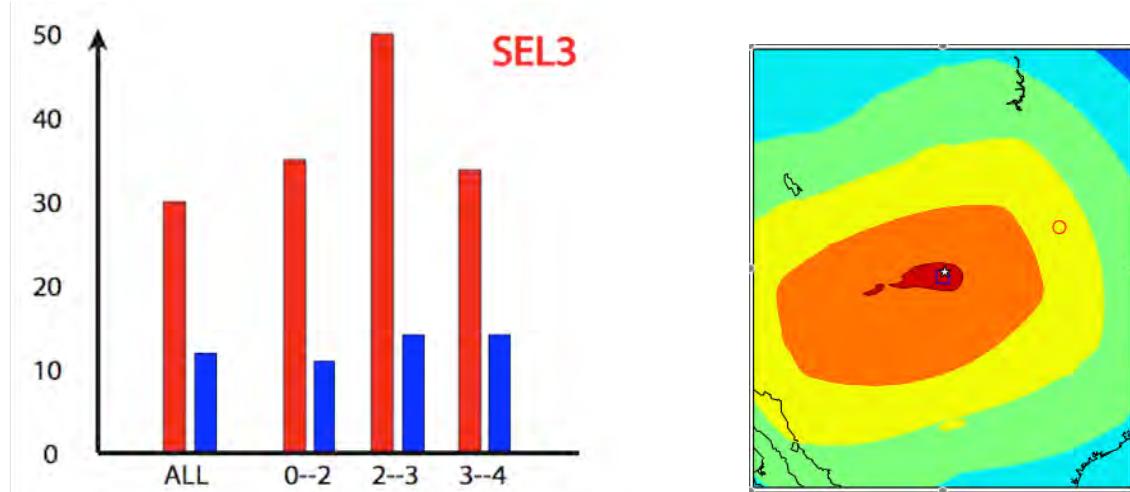
- A prior distribution  $P_T(e)$  over sets of events  $e$  for a given interval  $T$ . Each event is defined by time, location, depth, and magnitude. Event times are distributed according to a time-homogeneous Poisson process. (This is easily modified to allow for aftershocks.) Magnitudes are distributed according to the Gutenberg–Richter distribution. Locations for naturally occurring events follow a distribution estimated from historical data, whereas man-made events have a spatially uniform distribution with near zero depth.
- For each station and each true phase, a detection probability distribution is computed given the event magnitude, depth, and distance. The probability model (logistic regression with local input features) is calibrated from historical data that document which events were detected by which stations. (Figure 1(a)). Given that wave propagation is similar from a particular event region to each station for both earthquakes and explosions, there is good reason to use the numerous earthquakes to characterize the detection statistics of rare explosions.
- For each detection, the following predictive distributions for attributes of the detection (at that station) given attributes of the event:
  - A distribution for the predicted arrival time, given by the event time plus a travel-time distribution for the true phase from the event location to the station. NET-VISA currently uses the IASPEI91 1-D model with station-specific corrections, i.e., the same model as the IMS. Uncertainty is modeled as a Laplacian residual, as illustrated in Figure 1(b), that includes pick error.
  - A distribution for the measured amplitude given the event magnitude and distance – i.e., an attenuation model. As a starting point, the log amplitude is assumed to decay linearly with distance, with Gaussian error. This model ignores known deviations from simple geometrical spreading and attenuation models, which we expect to incorporate as the project progresses.

- A multinomial distribution for the measured phase label given the true phase label, again estimated from historical data.
- Distributions for the measured azimuth and slowness given the true azimuth and slowness for the event also follow a Laplacian distribution.
- Finally, for each station a time-homogeneous Poisson for spurious detections, each with its amplitude drawn from an empirically estimated distribution (a mixture of Gaussians), its phase label from an empirically estimated multinomial distribution, and its azimuth and slowness from uniform distributions. The assumption of uniformity for azimuth and slowness is questioned by Koper *et al.* (2010) and is valid only for a subset of stations in our data.



**Figure 1. (a) P-phase detection probability at ASAR for magnitude 3.5 events, as a function of distance. (b) Laplacian fit to ASAR travel-time residual with respect to iasp91 predictions with station-specific corrections. Note the bias of 0.6 seconds even after station-specific correction; our model automatically compensates for this bias.**

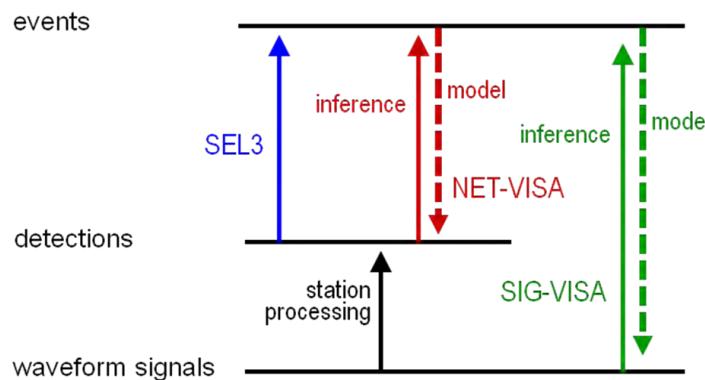
NET-VISA is currently being evaluated with a view to eventual deployment by the CTBTO. The evaluation compares the output of NET-VISA to the SEL3 bulletin (the final automated bulletin) produced by the CTBTO. For the purposes of this comparison, “ground truth” was defined by the LEB (Late Event Bulletin), which is generated by post-processing the SEL3 output by a team of expert analysts. The LEB itself is imperfect, but it provides a baseline of performance for a labor-intensive monitoring system. The distribution parameters for NET-VISA were determined based on 2.5 months of IMS data and tests were run on one week of held-out data. Details of the evaluation appear in the paper by Le Bras *et al.* (2011) in these proceedings. In summary, the low-magnitude (i.e., below 4) event-detection failure rate for NET-VISA is 2.5–3 times lower than that for SEL3, as shown in Figure 2(a). This improvement is achieved despite the fact that the “ground truth” LEB is derived from SEL3 and is therefore perhaps biased towards SEL3’s interpretation of the data. As illustrated in Figure 2(b), NET-VISA also finds events recorded by local networks (e.g., NEIC, JMA, PRU, NNC) that do not appear in LEB. Finally, NET-VISA was tested on one week of data that included the DPRK test of May 25, 2009; NET-VISA formed the event correctly with an accurate location based on associating 53 detections from the IMS, of which 50 were also associated in the LEB. (SEL3 included only 39 detections for the event.)



**Figure 2.** (a) Comparative detection failure rate for SEL3 and NET-VISA, for a test set of 852 events (1 week). Rates are shown for all events and for specific magnitude ranges. (b) Posterior location density calculated by NET-VISA for NEIC Event 12054557 (Colorado) using detections from three IMS stations (TXAR, PDAR, and ANMO), with the most likely location shown by the blue square. The white star shows the NEIC location using 19 stations; the red circle shows the nearest SEL3 location (approx. 10 degrees away); the event was not recorded in LEB.

### Signal-based Bayesian monitoring

Detection-based systems such as NET-VISA are fundamentally limited by the ability of the underlying signal detection algorithm. Our current project will develop the required statistical models and algorithms for a *signal-based Bayesian monitoring system*, which takes as its observed data the network-wide raw signal (waveform) data itself, thereby avoiding the need for hard-threshold local detection algorithms. Whereas NET-VISA has established the feasibility and credibility of Bayesian monitoring, SIG-VISA will be a qualitatively different and technically more sophisticated approach requiring fundamental advances in understanding. The basic relationship between NET-VISA and SIG-VISA is shown in Figure 3.



**Figure 3.** SEL3, NET-VISA, and SIG-VISA. SIG-VISA will make direct use of waveforms, whereas SEL3 and NET-VISA make use of a detection time series.

Signal-based Bayesian monitoring has several advantages over detection-based approaches:

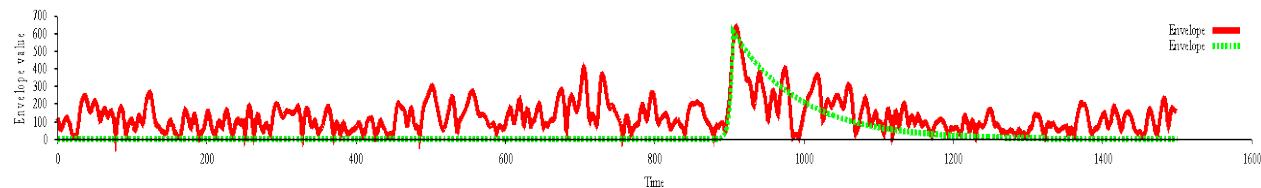
- Signal-based Bayesian monitoring is potentially more sensitive to low-SNR signals because it uses Bayesian inference to determine whether or not signal characteristics at *all* stations are consistent with the occurrence of an event at a given location and magnitude. Roughly speaking, signal-based Bayesian monitoring improves over detection-based methods in much the same way that array stations improve over single seismometers, by summing the signal recorded at multiple sensors to increase SNR.

- Signal-based Bayesian monitoring can also take advantage of correlations in waveform properties across multiple signals, as used in waveform matching algorithms such as double-differencing (Waldhauser and Ellsworth, 2000; Schaff and Richards, 2004).
- If signals from two events overlap at a given station, a detection-based system may record only a single, large detection, which may be incompatible with the locations and magnitudes of both events. In the signal-based approach, the overlapping signal is expected and may in some cases be separated into its components by the inference process (given time-separated signals from at least one other station). Stone et al. (1999) cite a similar advantage for their “unified tracking” approach.

We report here on two aspects of the problem: formulation of spatially correlated waveform models and connections between Bayesian and other “joint” methods such as double-differencing and joint hypocenter determination.

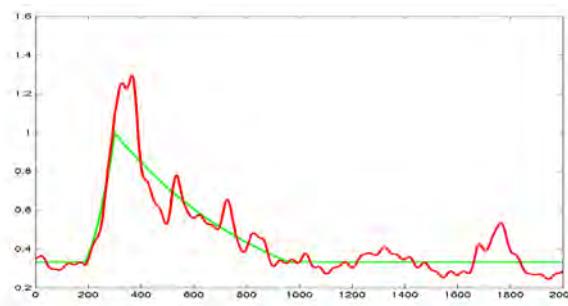
#### Waveform models: Preliminary proposal

The first research challenge lies in formulating a complete and consistent probability model that supports the computation of the likelihood of observing the measured signal traces, given any hypothesized collection of events. Fitting waveform envelopes is well established (e.g., Mayeda et al., 2004), and the simplest approach involves a low-dimensional parametric envelope descriptor (cf. Huseby et al., 1998) (e.g., triangular or paired-exponential) whose arrival time, amplitude, and spread depend on the distance and magnitude of the event; an example of a paired-exponential envelope template is shown in Figure 4.



**Figure 4.** A paired-exponential envelope template (green) superimposed on an actual waveform envelope (red).

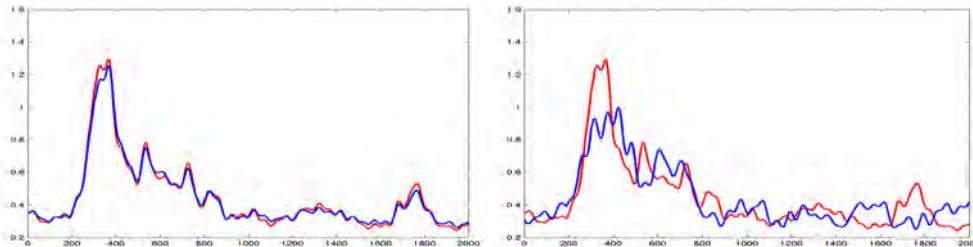
Actual envelopes are not well-modeled by “template plus i.i.d. noise”; typically, they exhibit more temporally correlated “macro-variation” resulting from path effects. Furthermore, envelope (or waveform) shape is highly repeatable across events with the same location and type – this is the basis for waveform cross-correlation. To achieve these properties we adopt a hybrid parametric/nonparametric model in which the realized envelope for a given event is the product of a parametric mean envelope (magnitude and distance dependent) and a stochastic modulation signal. A simple modulation model might be based on a random linear combination of Fourier basis functions (see Figure 5). Ultimately we will learn a model from historical data.



**Figure 5.** A sampled envelope (red) formed from a parametric template (green) multiplied by a random linear combination of Fourier basis functions.

A basic model of envelope variation would assign each event an independently sampled log-modulation signal, using a Gaussian prior on the basis coefficients. A nonparametric extension adopts a *Gaussian process* model for the

basis function parameters. This captures correlations among event envelopes that decay with distance (analogous to *correlation matching*); an illustrative example is shown in Figure 6.



**Figure 6:** (a) A reference envelope (red) and a sampled envelope (blue) from the Gaussian-process generative model for a nearby event location. (b) The same reference envelope (red) and a sampled envelope (blue) from a distant event location. Here, “nearby” and “distant” are relative to the Gaussian process correlation distance.

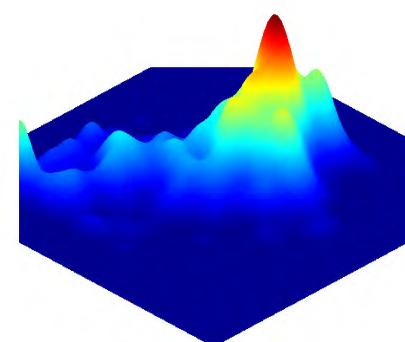
### Bayesian localization, joint hypocenter determination, and double-differencing

Accurately inverting event hypocenters is made difficult by velocity heterogeneity, which introduces correlations to the travel-time residuals. This issue has been addressed by the simultaneous inversion of hypocenters with station-specific corrections (e.g., Douglas, 1967), parameterized adjustments to the travel time prediction (Myers *et al.*, 2007, 2009, 2011), or with velocity structure (e.g., Thurber, 1983; Zhang & Thurber, 2003), and by difference-based methods such as double-differencing (Waldhauser and Ellsworth, 2000). The latter approach has the advantage of naturally integrating differential arrival time data obtained through waveform matching techniques such as cross-correlation; differences obtained in this way are more precise than those obtained by subtracting picked arrival times and so provide higher-quality location estimates (Schaff and Richards, 2004). We have developed a Bayesian method for the joint inversion of event hypocenters that accounts for correlated travel-time residuals as a fully integrated part of the probabilistic inference. We refer to this method as Bayesian double-differencing (BDD), because it resembles the method of double-differencing both in the sense of correcting for velocity heterogeneity and in its ability to exploit waveform-based relative arrival times. Unlike double-differencing and related techniques, we solve for absolute as well as relative locations, similar to Myers *et al.* (2007, 2009, 2011).

**Model:** Our approach models the slowness (inverse velocity) field as a Gaussian process (GP). In a GP, values at any pair of points are jointly Gaussian, with a covariance kernel that typically depends on the distance between the points. In standard GP regression (a.k.a. *kriging*), the posterior random field is inferred based on observations at a finite set of points; in our case, however, we do not observe the slowness values directly, and our goal is not to estimate the slowness field itself (although a posterior estimate is obtainable from the inference process). Instead, we make use of the fact that a GP slowness field induces a GP residual field at each station, with the covariance kernel of each residual field given by a double integral along event-station paths of the slowness field covariance kernel (Rodi and Myers, 2007). Because these station-specific residual fields are all induced by the same underlying GP slowness field, the travel-time residuals themselves are jointly Gaussian across stations. (Note that unlike difference methods, which consider the relationships between event residuals independently at each station, our model shares information naturally between nearby stations.)

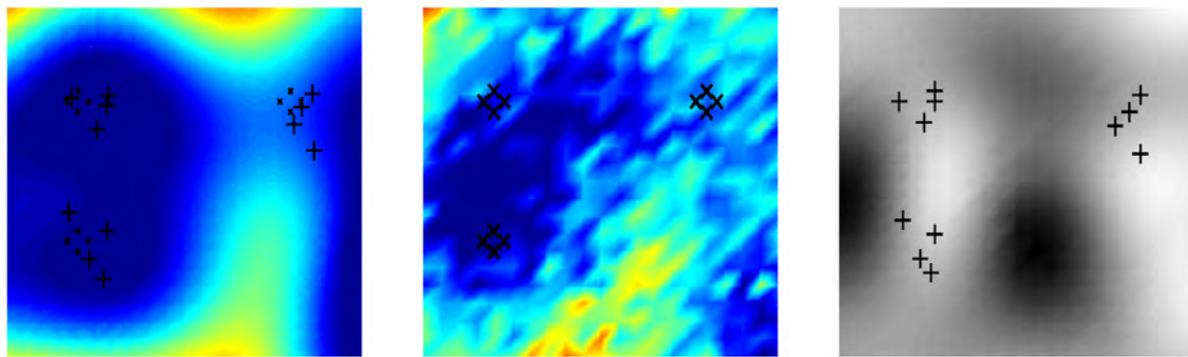
The joint distribution on the residuals given the event locations constitutes a forward model, which we combine with a prior distribution on event locations and invert using Bayes’ rule to yield a posterior distribution on locations conditioned on arrival times. Of course, the true arrival times are not observed; we must extend the model to reflect sensor noise in the form of “pick errors.” This is accomplished by an additional Gaussian noise term that is added independently for each observation as part of the forward model.

It is similarly possible to incorporate arrival-time differences obtained by waveform matching: these differences are represented in the forward model as the differences in the true arrival times, plus a



**Figure 7:** Example posterior marginal distribution for a single event location.

Gaussian noise term (again independent for each difference observation) with lower variance than that of the pick error to reflect the higher precision attained by waveform matching. The Bayesian inversion seamlessly makes use of both the absolute time constraints provided by picked arrival times along with the relative time constraints of the difference observations when available; the lower variance ascribed to difference observations causes them to be treated as stronger constraints, analogously to the higher weighting commonly given to these constraints in double-difference methods (also see Myers et al., 2011, these proceedings).

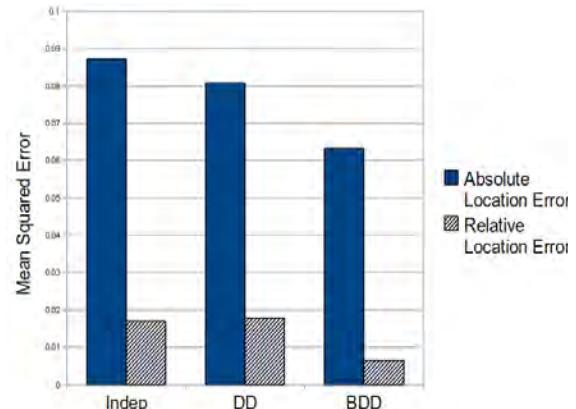


**Figure 8.** (a) Simulated 2-D event clusters located using BDD (+) along with their true locations (x). Stations are positioned at the four corners and each event is observed by three of the four stations. The background shows the MLE slowness field given the inferred locations. (b) The true slowness field, with true event locations marked. (c) Variance of the slowness field posterior (darker is higher).

**Inference:** The posterior distribution over the event locations themselves has no simple analytic form (Figure 7 shows an example), so the inversion is nontrivial to perform. We instead apply Markov Chain Monte Carlo (MCMC), a standard technique for Bayesian inference, to sample from the posterior and estimate the expected locations of the joint set of events (e.g. Lomax et al., 2000; Myers et al., 2007). This is time-consuming but tractable for the small-scale simulations we have considered thus far.

Although the goal of the inference is to produce accurate locations, the model also defines a posterior distribution over the slowness field, which is simple to calculate given a point estimate of the event locations. Figures 8(a) and 8(b) compare a slowness field inferred from 12 events to the true slowness field, showing that our method effectively recovers the tomography while inverting for event locations. In this sense it resembles other tomographic approaches (e.g. Thurber, 1983; Zhang & Thurber, 2003), but maintains a full posterior distribution over the tomography instead of a point estimate, similar to Shapiro and Ritzwoller (2004), Pasyanos et al., (2006) and Hauser et al., (2011). As illustration, Figure 8(c) shows the variance of the slowness field posterior; note that we are most confident of the tomography along ray paths between events and nearby stations (located at the four corners).

**Experiments:** We have conducted initial simulation experiments with clustered events detected by three of four stations and arrival times sampled from the Gaussian process prior. These experiments did not make use of waveform matching. Several distance-weighting strategies for double-differencing were evaluated, including the one described by Waldhauser & Ellsworth (2000); the results presented in Figure 9 are from the best-performing strategy, which sets the weight of each difference residual constraint equal to the prior covariance of the two corresponding residuals in the Bayesian model. As Figure 9 shows, the Bayesian method significantly outperformed both double-differencing as well as a simple independent inversion of the event locations, both in the absolute location error and in the relative locations between events.



**Figure 9:** MSE in absolute and relative locations (22 simulation runs, 5 events per run).

## **CONCLUSIONS AND RECOMMENDATIONS**

Prior results suggest that Bayesian monitoring is a promising technique for analyzing streams of parametric detections from multiple stations to form a global event bulletin and may be preferable to existing deployed methods for global association. Our current work, in its very early stages, is aimed at extending Bayesian monitoring with generative models of waveform envelopes to improve detection and association and performing joint inference of event locations and slowness fields to improve localization accuracy.

## **ACKNOWLEDGEMENTS**

We gratefully acknowledge the support of DTRA for this work; also the support of CTBTO for the work on NET-VISA and the technical cooperation of CTBTO personnel including Ronan Le Bras and Jeff Given. We thank Sheila Vaidya of LLNL and Ola Dahlman, former Chair of CTBTO Working Group B, for their moral support and encouragement in this project.

## **REFERENCES**

Antolik, M., G. Ekstrom, and A. Dziewonski (2001). Global event location with full and sparse datasets using three-dimensional models of mantle  $P$ -wave velocity, *Pure Appl. Geophys.* 158: 291–317.

Arora, N. S., S. J. Russell, P. Kidwell, and E. Sudderth (2011b). Global seismic monitoring: A Bayesian approach. In *Proc. AAAI-11*, San Francisco.

Arora, N. S., S. J. Russell, P. Kidwell, and E. Sudderth (2011a). Global seismic monitoring as probabilistic inference. In *Advances in Neural Information Processing Systems 23*, MIT Press.

Arora, N., Russell, S., and Sudderth, E. (2009). Vertically Integrated Seismic Analysis II: Inference. *Eos Transactions of the American Geophysical Union*, 90(52), Fall Meeting Supplement, Abstract S31B-1713.

Arora, N., Russell, S., de Salvo Braz, R., and Sudderth, E. (2010b). Gibbs sampling in open-universe stochastic languages. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI 2010)*, Catalina Island, California.

Bell, K. L., Ephraim, Y., and Van Trees, H. L. (2000). A Bayesian Approach to Robust Adaptive Beamforming. *IEEE Transactions on Signal Processing*, 48(2),

Candy, J. V. (2009). *Bayesian Signal Processing: Classical, Modern, and Particle Filtering Methods*. Wiley Interscience.

Dodge, D.A., and K.F. Sprenke (1992). Improvements in mining induced microseismic source locations at the Lucky Friday mining using an automated whole-wave-form analysis system, *Pure Appl. Geophys.* 139: 609–626.

Douglas, A. (1967). Joint epicentre determination. *Nature* 215: 47–48.

Duijndam, A. J. W. (1988a). Bayesian estimation in seismic inversion—Part I: Principles. *Geophys. Prosp.*, 36: 878–898.

Duijndam, A. J. W. (1988b). Bayesian estimation in seismic inversion—Part II: Uncertainty analysis. *Geophys. Prosp.*, 36: 899–918.

Flanagan, M.F., S.C. Myers, K.D. Koper (2007). Regional travel-time uncertainty and seismic location improvement using a 3Dimensional *a priori* velocity model, *Bull. Seismol. Soc. Am.* 97: 804–825.

Gibbons, S. J. and Ringdal, F. (2006). The detection of low magnitude seismic events using array-based waveform correlation. *Geophys. J. Int.* 165, 149 – 166.

Hauser, J., K.M. Dyer, M.E. Pasmanos, H. Bungum, J.I., Faleide, S.A. Clark, and J. Schweitzer (2011). A probabilistic seismic model for the European Arctic, *J. Geophys. Res.* 116: B01303, doi:10.1029/2010JB007889

Harris, D.B (1991). A wave-form correlation method for identifying quarry explosions, *Bull. Seismol. Soc. Am.* 81; 2395–2418.

Koper, K. D., Seats, K., and Benz, H. (2010). On the composition of Earth's short-period seismic noise field. *Bull. Seismol. Soc. Am.* 100(2); 606–617.

Lomax, A., J. Virieux, P. Volant and C. Berge (2000). Probabilistic earthquake location in 3D and layered models: Introduction of a Metropolis-Gibbs method and comparison with linear locations, in *Advances in Seismic Event Location* Thurber, C.H., and N. Rabinowitz (eds.), Kluwer, Amsterdam, 101–134.

Mayeda, K., Hofstetter A., O'Boyle J. L., Walter, W. R. (2003), Stable and transportable regional magnitudes based on coda-derived moment-rate spectra. *Bull. Seismol. Soc. Am.* 93: 224–239.

Myers, S.C., G. Johannesson, and W. Hanley (2007). A Bayesian hierarchical method for multiple-event seismic location. *Geophys. J. Int.* 171: 1049–1063.

Myers, S.C., G. Johannesson, and W. Hanley (2009). Incorporation of probabilistic seismic phase labels into a Bayesian multiple-event seismic locator. *Geophys. J. Int.* 177: 193–204.

Myers, S.C., G. Johannesson, and N.A. Simmons (2011). Global-scale P-wave tomography optimized for prediction of teleseismic and regional travel times for Middle East events: 1. Data set Development, *Jour. Geophys. Res.*, 116, B04304, doi:10.1029/2010JB007967.

Myers, S.C., G. Johannesson, D.A. Dodge, and N.A. Simmons (2011), Combining analyst and waveform-correlation-based arrival time measurements in a Bayesian multiple-event location algorithm, these Proceedings.

Pasyanos, M. E., G. A. Franz, and A. L. Ramirez (2006), Reconciling a geophysical model to data using a Markov chain Monte Carlo algorithm: An application to the Yellow Sea–Korean Peninsula region, *J. Geophys. Res.*, 111, B03313, doi:10.1029/2005JB003851.

Peterson, J. (1993). *Observation and Modeling of Seismic Background Noise*. USGS Open-File Report 93-322.

Rodi, W. L., and S. C. Myers (2007), Modeling travel-time correlations based on seismicity kernels and correlated velocity anomalies, in *Proceedings of the 29th Monitoring Research Review: Ground-Based Nuclear Explosion Monitoring Technologies*, LA-UR-07-5613, Vol. 1, pp. 463–471.

Russell, S., Vaidya, S., and Le Bras, R. (2010). Machine Learning for Comprehensive Nuclear-Test-Ban Treaty Monitoring. *CTBTO Spectrum*, 14: 32–35.

Russell, S., Arora, N., Jordan, M., and Sudderth, E. (2009). Vertically Integrated Seismological Analysis I: Modeling. *Eos Transactions of the American Geophysical Union*, 90(52), Abstract S33D-08.

Schaff D.P., and P.G. Richards (2004). Lg-wave cross correlation and double-difference location: Application to the 1999 Xiuyan, China, sequence, *Bull. Seismol. Soc. Am.* 94: 867–879.

Schultz, C.A., S.C. Myers, J. Hipp, and C. Young (1998). Nonstationary Bayesian kriging: a predictive technique to generate corrections for detection, location, and discrimination, *Bull. Seismol. Soc. Am.*, 88: 1275–1288.

Shapiro, N. M., and M. H. Ritzwoller (2002). Monte-Carlo inversion for a global shear velocity model of the crust and upper mantle, *Geophys. J. Int.* 151: 88–105.

Simmons, N.A., A.M. Forte, L. Boschi and S.P. Grand (2010a), GyPSuM: A joint tomographic model of mantle density and seismic wave speeds, *J. Geophys. Res.* (under review).

Simmons, N.A., S.C. Myers, G. Johannesson (2011). Global-scale P-wave tomography optimized for prediction of teleseismic and regional travel times for Middle East events: 1. Tomographic inversion. *Jour. Geophys. Res.*, 116, B04305, doi:10.1029/2010JB007969, 2011.

Stone, L. D., Corwin, T. L., and Barlow, C. A. (1999). *Bayesian Multiple Target Tracking*. Artech House.

Taylor, S. R., Velasco, A., Yang, X., Maceria, M., and Phillips, W. S. (2001). Bayesian tomography applied to seismic event identification problems, in *Proceedings of the 23<sup>rd</sup> Seismic Research Review: Worldwide Monitoring of Nuclear Explosions*, LA-UR-01-4454, Vol. 1, pp. 440–447.

Thurber, C. H. (1983). Earthquake locations and three-dimensional crustal structure in the Coyote Lake area, central California. *J. Geophys. Res.* 88(B10), 8226–8236. American Geophysical Union.

Waldhauser, F. and Ellsworth, W. L. (2000). A Double-Difference Earthquake Location Algorithm: Method and Application to the Northern Hayward Fault, California. *Bull. Seismol. Soc. Am.* 90(6): 1353–1368.

Yang, X., Bondar, I., Bhattacharyya, J., Ritzwoller, M., Shapiro, N., Antolik, M., Ekström, G., Israelsson, H., and McLaughlin, K. (2004). Validation of Regional and Teleseismic Travel-Time Models by Relocating Ground-Truth Events. *Bull. Seismol. Soc. Am.* 94: 897–919.

Zhang, H., & Thurber, C. H. (2003). Double-difference tomography: The method and its application to the Hayward fault, California. *Bull. Seismol. Soc. Am.* 93(5): 1875.